PC Support Advisor

The Essential Resource for PC Support Professionals

Update 63

On This Month's Utility Disk

Latest version of Telix communications program.

Fully-functional virus scanner.

"*Visually Superior*" - says Robert Schifreen.

The Windows bandwagon continues to roll. Microsoft has just announced that it has sold a total of 40 million copies of Windows since the launch of 3.0. While much of this is down to the fact that Windows is a genuinely useful product for many people, clever marketing also has a lot to do with it. I'm reliably informed that Microsoft no longer sells plain MS-DOS to OEMs - they have to take Windows too. And there's no point in buying Windows and not shipping it, which is why so many machines come bundled with Windows.

Contrast this with IBM. Only a handful of machines come pre-installed with OS/2, and this handful doesn't include any of IBM's own PCs. Yup - if you buy an IBM PC you'll almost certainly get DOS and Windows with it. If IBM can't even convince itself that OS/2 is better than Windows, how can it possibly convince its customers?

Whatever the operating system, it's the applications that make the whole machine useful. Developing useful applications is no easy task, of course, and many companies strive to help developers and programmers have an easier life.

This month, the award for making my own life easier goes to Microsoft for Visual Basic. Not the Windows version, but the DOS product. I've spent the last half-dozen weekends working on an application for my own use in order to help me manage the production of PC Support Advisor as well as the monthly index files, the Utility Disks and the Briefing Tapes. In about 12 days of programming I've managed to put together something that would have taken me months under any other language or environment, with the possible exception of Borland's Pascal 7. But I prefer Basic to Pascal because I know the language more fluently and I have a job to do - I don't have time to re-learn Pascal.

I'm also highly impressed with the speed with which I can make changes to a VBDOS application. Enlarging a list box or moving a window takes half a dozen mouse clicks and the work is done - there's no real code to alter at all. I'm not alone in my admiration for VBDOS - we regularly receive comments about it in the PCSA office. There are complaints too, though these are few in number and mainly concern the speed at which the development environment runs out of working RAM. Apart from that, though, I find it one of the most productive environments I've ever used. And the fact that I can produce a real Windows version of my application in 10 minutes, just by recompiling under VB for Windows, makes me very happy.

It's not surprising that, according to Microsoft, Visual Basic is outselling all other languages, including C and C++, by a factor of five or more. When you have users banging on your door to solve their problems, you need to be able to sort out their troubles quickly and with as little fuss as possible. This often means writing a small program that runs alone, or which talks to an existing application. For my money, VB is the best way to do these types of jobs.

PCSA

PC Support *Advisor*

PC Support *Advisor* is a publication of **International Technology Publishing**.

ISSN: 1031-3966

Editor: Robert Schifreen

Editor in Chief: Les Bell

Publisher: Chris Carvan

Email: hex@cix.compulink.co.uk

International Technology Publishing Inc.
North American Subscriptions
Suite 200
555 De Haro Street
San Fransisco, CA94107
USA

Telephone: (415) 255 1295
Facsimile: (415) 255 8496

International Technology Publishing Inc.
European Subscriptions
13 Callcott Street
London
W8 7SU
United Kingdom

Telephone: (071) 724 9306
Facsimile: (071) 706 3296

International Technology Publishing Pty Ltd.
SE Asian Subscriptions
138 Cecil Street
Cecil Court #13-03
Singapore
0106

Telephone: (65) 227 4187
Facsimile: (65) 227 2885

International Technology Publishing Pty Ltd.
Australasian Subscriptions
Suite 511
185 Elizabeth Street
Sydney, NSW, 2000
Australia

Telephone: (02) 261 4683
Facsimile:(02) 261 4741

Understanding PC Display Adapters

First there was MDA, then came CGA, EGA, VGA, SVGA, MCGA, XGA and more. What does it all mean? Julian Moss has the details.

The display adapter is the hardware device that turns ASCII text or a memory representation of a graphical image into the set of signals that enable a monitor to display it on the screen. As interest in high resolution graphics has increased, great changes have been made in display adapter technology. PCs are now capable of sparkling graphics performance undreamt of just a couple of years ago.

This article looks at how display adapters work, examines the aspects of their design that affect their capabilities and their performance, and provides guidance on how to go about selecting display adapters for use within your organisation.

Display Basics

In a conventional computer display such as a mainframe terminal, data is typically sent from the computer to the display over a serial communications link using some kind of protocol. The protocol includes commands for things like positioning the cursor and setting the colour of text, as well as commands that say "here is a string of characters to be displayed" and "here is a sequence of numbers representing graphics." This method of driving a display has advantages where the terminal is remote from the computer, but it is slow, being limited by the speed of the link.

PC displays are memory mapped. This means that there is an area in memory, directly addressable by the processor, which contains a representation of what is displayed on the screen. This memory is also accessible to the display adapter hardware. If the processor makes a change to something stored in display memory, the change appears almost instantly on the screen. Memory mapped displays are fast and flexible, but are only possible where there is a close relationship between the computer and the display hardware, as is the case in a PC.

Frame Buffer

The area of memory used to store the display information is called the frame buffer. In the design of the IBM PC, certain areas of memory have been reserved for use as frame buffers by different types of display adapter. An area starting at address B0000h (or B000:0 in segmented terminology) is used for monochrome displays in text mode; colour displays use an area beginning at B8000h. For graphics displays the buffer starts at address A0000h.

In text mode, two bytes of memory are used for each character displayed. The first byte is the code of the character itself. The second defines its attributes: the foreground and background colours, the foreground intensity and whether or not the character is blinking. A screen of 25 rows of 80 characters therefore needs 4,000 bytes of memory.

Even the most basic PC display adapter has 16 KB of memory so several separate "pages" of text can be supported, though the number of software packages that take advantage of this facility is small. By default, page 0 is normally displayed on screen, but a simple command can put the contents of page 1 on screen instead. This allows you to build up a picture gradually, in the background, perhaps by performing a number of complex calculations, then flash it instantly onto the screen. Animation can also be achieved by loading a succession of frames into pages and flicking between them.

Character Bytes

The character byte contains the ASCII code of the character in the case of values 0 to 127, and a code representing a special IBM character for the remaining 128. These extra 128 characters include accented letters as well as line and block graphics. The video controller uses the character byte to access a look-up table - the character table - which holds a bitmapped representation of that character. This tells the controller how the character should be displayed on screen.

The standard, default character table is stored in the display adapter ROM. The IBM Color Graphics Adapter (CGA) - the earliest PC colour display - had a resolution of 640 x 200 and used a chunky 8x8 character box to display 25 rows of 80 characters. The Enhanced Graphics Adapter (EGA) increased the resolution to 640 x 350, and used an 8x14 box for better shaped characters, though 8x10 and 8x8 characters were also available to give 35 or 43 lines on the screen. The Video Graphics Array (VGA) adapter uses a 720 x 400 resolution and a 9 x 16 character box in standard text mode. Up to 50 rows can be displayed using smaller fonts.

Fonts

Both EGA and VGA displays support the use of user-definable fonts, which allow you to load your own character table into the video adapter's memory. This feature is often used by software such as the Symantec's Norton Utilities to create a graphical look

"Third-party display adapters with extra features, such as Super VGA adapters, have their own unique additions to the set of video BIOS functions called extended functions."

in text mode. Though user-definable fonts are a standard feature supported by the EGA/VGA BIOS, they can sometimes cause problems. Task-switching utilities, for example, may not automatically restore the user-defined font when the screen is switched from one application to another and back again.

The system BIOS provides functions for text output, cursor positioning, scrolling and clearing the screen. However, the BIOS functions are primitive and slow. Most text mode programs write characters and attributes directly into the frame buffer, only using the BIOS video functions for things like controlling the cursor and scrolling the screen.

In graphics mode, the data in the frame buffer memory represents the attributes of the individual pixels that make up the picture: the term "pixel" means picture element. The characteristics of each pixel can be individually set. IBM called this "all points addressable" (APA) graphics, to distinguish it from the character-oriented block graphics available in text mode. BIOS functions make it possible to read and write the value of an individual pixel, but once again using the BIOS to display graphics is slow. Bypassing the BIOS is more difficult, though, due to the different ways in which the graphics screen can be represented in the frame buffer.

Planes

At lower colour depths, the memory is often arranged as a number of planes, each representing a colour or intensity level, and each using 1-bit per pixel. In other modes, the memory is arranged more as you might expect, as a two dimensional matrix with one or more bytes per pixel. Unless you're writing a graphics program, there's no reason to worry about this. However, the lack of hardware support for anything more sophisticated than pixel-plotting is one reason why the PC has for a long time been regarded as a poor machine for graphics.

Adapter Compatibility

BIOS support for the PC display is provided by the Interrupt 10h set of functions. These include support for setting the display mode (the number of rows, columns and colours, text or graphics), the shape and position of the cursor, the display page, reading and writing characters or pixels and scrolling. These functions are standard for all PC display adapters and in the case of the CGA and monochrome adapters used in the original PC they are part of the system BIOS.

The EGA and subsequent display adapters introduced additional BIOS video functions, such as to set the border colour and palette registers. They also introduced a function for changing the use of the text mode blink attribute to specify bright background colours instead - in earlier display adapters this had to be done by writing to one of the controller's registers.

Because of the significant hardware differences between EGA adapters and above and their predecessors, these display cards provide their own replacement set of Interrupt 10h BIOS functions in a ROM mounted on the card. The standard memory location for this ROM is addresses C0000h - C7FFFh. However, some PCs with integrated graphics adapters locate it at E0000h so that all the ROM code can be stored on a single 128 KB chip.

Additional Features

Third-party display adapters with extra features, such as Super VGA adapters, have their own unique additions to the set of video BIOS functions, called extended functions. Since they are not part of the notional IBM PC standard, these functions can only be used by software written specifically for that display adapter.

Controller Chips

Early PC display adapters used a controller chip type 6845. This chip has a set of registers which can be used to change the basic characteristics of the display. Later display adapters add to this set of registers, whilst emulating the features of the 6845 for the benefit of software written for the older display standards.

Most of the time, software should have no need to access the video controller registers: BIOS functions are provided for most of the things that programs would want to do. However, because the register functions are documented, programmers have manipulated them directly, either for speed, or to do things that IBM never intended such as smooth scrolling. Consequently, it is important that display adapters provide register-level compatibility as well as BIOS-level compatibility with the IBM VGA and earlier standards.

Signals

The signals output by a display adapter consist in essence of a set of pulses which determine whether the electron beams in the CRT monitor should be bright, off or something in between. The frequency of these pulses is related to the number of pixels the beam must scan every second. This frequency in turn is derived from the resolution of the picture, and the refresh rate - the number of times the picture must be re-displayed every second.

This latter figure has become in-

PC Display Adapters

creasingly important as concern has been voiced over the effects on the user of flicker. The higher the refresh rate, the less noticeable the flicker; a rate of 72 Hz (Hertz, where 1 Hz is once per second) is currently considered ergonomically acceptable. Standard VGA adapters use a refresh rate of 60 Hz at 640 x 480 resolution, whilst lower rates still are commonly used for Super VGA. In an interlaced display, two passes of the screen are made to try to reduce the objectionable flicker that would be caused by a refresh rate of 43.5 Hz. Sometimes this can work well, but usually you get an equally annoying shimmer instead.

Pulses

A vertical sync pulse - equal to the refresh rate - tells the monitor when a scan is about to begin. Another signal, the horizontal sync pulse, marks the start of each scan line. The frequency of this signal is related to the refresh rate and the vertical resolution - the number of rows of pixels - in the picture. You can't calculate the frequencies by simple multiplication, though, because they must also take into account the time lost when the beam moves back from the end of one scan line to the start of the next (horizontal retrace) and from the bottom of the screen to the top (vertical retrace). The pixel pulse frequency is called the dot clock.

Frequencies vary from a dot clock of 25.175 MHz and a horizontal scan rate of 31.5 KHz for a 640 x 480 display at 60 Hz, to a dot clock of 77 MHz and a scan rate of 58 KHz at 1024 x 768 at 72 Hz. At 1280 x 1024 you need a dot clock of 110 MHz and a scan rate of 64 KHz to get even a 60 Hz refresh rate. Not surprisingly, the hardware design becomes more demanding and hence more expensive at these higher rates, which is one reason why the maximum supported refresh rate drops off as resolution increases.

Multisync

Older monitors could only lock on to sync pulses at fixed frequencies: the 31.5 KHz scan rate used by standard VGA monitors also allowed a 70 Hz

Graphics Accelerator Features Explained

Graphics accelerators work by incorporating hardware dedicated to performing a particular graphics function. Some of the terms used to describe the functions they perform can seem a little opaque. Here is a description of some of the most common features:

Bit Block Transfer

Often abbreviated to BitBlt, a bit block transfer is used to move a block of pixels from one location to another. Without this feature, a pixel would have to be read from one location and then written back to another, passing across the data bus twice. With it, the operation can be performed within the display adapter, speeding up operations like scrolling.

Font Caching

Instead of writing the pixel pattern representing a character in a particular font each time it appears, the pattern is stored in the graphics card's memory so only the character code and its position need be sent to the card for subsequent instances of the character.

Hardware Clipping

In a windowing environment, each window is treated as a virtual screen. Programs typically write to the window without worrying about whether the information will be visible or not. The windowing software decides if the data should actually be displayed on the screen. Using hardware clipping, the graphics card, rather than the software, makes this decision. This saves the CPU some work.

Hardware Cursor

An image of the cursor can be loaded into the graphics chip, and this will then be displayed simply by sending its coordinates. Without this feature, the processor must take responsibility for drawing the cursor and saving and restoring the contents of the area over which it passes.

Hardware Windows

Instead of using the frame buffer memory as a direct representation of the screen image, the memory is divided up into smaller regions corresponding to individual windows. This speeds up many window operations by simplifying many of the computations necessary to determine where something should appear on screen.

Line Drawing

If this feature is present, the graphics adapter will draw a line of specified width between two given points. Without it, the software must work out where the line will go and plot the relevant pixels.

Polygon Fills

Using this feature, the graphics card can be given the boundary of a rectangle or other polygon, and told to fill the enclosed area with pixels of a specified colour. Without it, the processor would have to set each pixel individually.

Pattern Fills

Similar to the above, this tells the graphics adapter to fill an area with a predefined bit pattern. This is particularly useful for dithered colours in Windows.

"One of the problems users can encounter when using different display modes is that the size and position of the on-screen image changes between each one. This problem is caused by differences in signal timings."

refresh rate to be used in text mode, where the pixel resolution is 720 x 400. Today, display adapters can produce such a range of different modes that the use of multisync monitors, which can lock on to anything within a given range, is almost mandatory.

Shifting Images

One of the problems users can encounter when using different display modes is that the size and position of the on-screen image changes between each one. This problem is caused by differences in signal timings.

When a horizontal sync pulse is received, the CRT's electron beam is pulled back to somewhere off the left hand edge of the screen, and immediately begins to traverse across to the right. There is a short interval, called the front porch, before the pixels of that row are painted. After the row is finished there is a second interval called the back porch, after which the next sync pulse is received. A similar situation occurs at the start and end of each frame: the electron beams are turned off during the retrace for a period called the vertical interval.

In a fixed frequency monitor, the timings of these signals are standard and it is easy for the monitor to be set up at the factory so that the picture fits the displayable area of the CRT. However, the variations between different display modes and different display cards means that a multisync monitor has to adjust its gain - the speed with which the beam traverses the screen - dynamically. Some are better at doing this than others. Some display adapters, such as those based on the ET4000AX controller, come with a memory resident utility that can adjust the picture positioning by varying the timings of the signals produced by the card. However, the range of adjustment is small. This problem is really a monitor problem, and the only way to avoid it is to test the graphics card with a sample of each type of monitor you will be using before making a purchase.

Colour

Colours are created on screen by addition of the primary colours: red, green and blue. Before VGA, monitors were digital: each primary colour could be bright, dim or off. This gave a choice of 16 colours made up of dim and bright versions of the eight possible combinations of three colours. The number of bits per pixel or the number of colours are known as the colour depth. 16 colours can be represented conveniently in 4-bits, so two pixels can be stored in each byte. Where there is a direct correspondence between each pixel's value and its colour, the representation is known as direct mapped.

EGA's Enhanced Colour Display had two control signals for each colour, giving four brightness levels (including off) for each one. This allowed a maximum of 64 colours to be displayed. EGA kept the memory-efficient 4-bits per pixel scheme, but instead of direct mapping the 4-bits are used to select 16 colours out of a palette of 64. A default system palette let you use the display as if it was direct mapped, for those not requiring the additional colours.

With the advent of VGA came the use of analogue monitors. Having continuously variable levels for each of the three primary colours, a VGA monitor can display an almost infinite range of colours. Because the monitor must be controlled by a digital device, of course, the range of colours you can use is in practice not infinite. The maximum is determined by the resolution of the digital-to-analogue converter (DAC) which is built into the display adapter.

In a standard VGA card, the DAC has a six-bit resolution. Multiply that by three - one for each of red, green and blue - and you have 18-bits or a total of 262,144 different colours. VGA supports 256 simultaneous on-screen colours using one byte per pixel. A 256 entry colour look up table (CLUT) built into the DAC is used to convert between the pixel value and the actual colour to be displayed.

Other systems exist which can display even more on-screen colours for even greater photorealism. The TrueVision TARGA system is direct-mapped with 5-bits for each of the primary colours, giving 32,768 hues in total. This system is also known as HiColor.

IBM's XGA offers a direct-mapped system with 16-bits per pixel, made up

Resolution	**Colours**			
	16	**256**	**65 K**	**16 M**
640 x 480	256 KB	512 KB	1 MB	1 MB
800 x 600	256 KB	512 KB	1 MB	1.5 MB
1024 x 768	512 KB	1 MB	1.5 MB	2.5 MB
1280 x 1024	1 MB	1.5 MB	2.5 MB	4 MB
1600 x 1200	1 MB	2 MB	4 MB	6 MB

Table 1 - Adapter Memory by screen resolution

PC Display Adapters

of five each for red and blue and six for green. XGA uses the same 6-bit DAC as VGA. However, 6-bits per colour would require a rather inconvenient 18-bits per pixel, while with 5-bits per colour there would be an unused bit in every pair of bytes. The eye is more sensitive to changes in the level of green light than red or blue, so the extra bit is used for green. The result is a total of 65,536 simultaneous on-screen colours.

The ultimate as far as colour depth is concerned is the 24-bit colour or True Colour display. This direct mapped display mode requires a 3 x 8-bit DAC. It uses three bytes for every pixel, so this is the most demanding mode of all for video memory. A maximum of 16 million colours - more than could be discerned by the eye let alone displayed simultaneously on screen - is supported by a 24-bit display adapter.

Though all VGA and Super VGA graphics cards support 16-colour and 256-colour display modes, the ability to use higher colour depths varies from one model to another. However, most now support HiColor and XGA 16-bit colour modes, and may even offer True Colour support if an optional 24-bit DAC is installed in the card. In practice, though, the use of higher colour depths with ordinary graphics cards may be hampered by two things: lack of frame buffer RAM and poor performance.

"Most graphics chip sets are capable of supporting higher resolutions requiring at least 2 MB of video memory. If this is not installed as standard then sockets for the extra memory chips are generally provided on the board."

Moving Pictures

This article has concentrated on the display of still images on a PC monitor. In theory, there's no reason why a PC can't display moving pictures too. In the past, a number of companies have produced systems for capturing video from a VCR and playing it back on a PC's screen as a series of rapidly-changing images. Autodesk Animator is another popular way to produce moving pictures.

Today's most common solution is Microsoft Video for Windows, which scores highly in usability because it needs no extra hardware to play back movie clips on a standard PC under Windows. However, the processor speed of the PC plays a crucial part in ensuring that the image is of reasonable quality, as each frame has to be fetched and displayed in order to produce what appears to be a moving image. Play a .AVI (as used by Video for Windows) file at standard VGA resolution in a 2-inch window and it looks acceptable. Resize it to full screen and, in order to keep up with the sound track, the system will probably only show you every third or fourth frame.

Moving pictures take up a lot of disk space. At 25 frames per second, multiply the figures in Table 1 by 25 to see just how much space is needed to store just one second's worth of movie. Compression is a possibility, but if the decompression takes too long then it will affect performance.

Fractal compression looks like being the key to producing moving pictures that don't require hundreds of MB of hard disk space. The fractal compression algorithm developed by Iterated Systems can reduce a frame to literally 1 or 2% of its original size, and the company has demonstrated some truly remarkable software which plays back full-screen moving video (albeit under DOS rather than Windows).

Microsoft's vision for the future is to have moving video on every PC, so that finding out what's happening in the office next door is no more difficult than looking up a figure in a spreadsheet.

Video Memory

High resolution graphic images need a lot of storage. A standard VGA resolution (640 x 480) graphics display contains 307,200 pixels. This number more than doubles if you go to 1024 x 768. Increasing the colour depth to 8, 16 or 24-bits multiplies it a further two, four or six times.

Table 1 shows the rounded-up memory requirements for all of the popular VGA and Super VGA screen resolutions, at each colour depth. Most low to mid-price display adapters today come equipped with 1 MB of memory, which is adequate for the display needs of most users.

Most graphics chip sets are capable of supporting higher resolutions requiring at least 2 MB of video memory. If this is not installed as standard then sockets for the extra memory chips are generally provided on the board. It is worth checking this, though, because for today's price-conscious market some OEMs produce video cards without these extra sockets - even though a space is provided for them on the card - and if you get one you won't be able to upgrade it.

If you're considering a high resolution display adapter supporting 16-bit or 24-bit colour for use with image editing software then not only will you need sufficient memory on the card itself, you'll probably need an equival-

ent amount of extra main RAM over and above that needed to run the graphics software. Software that works with images usually keeps copies of the images in memory. Windows software will be able to use virtual memory created by the Windows swap file, but if it does, performance will be very slow.

Whilst Table 1 is a useful guide when checking the memory requirements of a video card, it can be wrong. If you want to use high refresh rates to reduce flicker as well as high resolutions and colour depths you may need to install double the memory that the resolution would appear to demand. To understand why, we need to look a bit more closely at how the memory in a display adapter is used.

Display memory is accessed by two things: the processor, which writes information to it whenever the display needs updating, and the video controller, which is constantly accessing every byte in the frame buffer as it generates the signals that cause the monitor to paint the image on the screen. A problem occurs if both the processor and the video controller try to access the same memory location at the same time - it isn't physically possible. To overcome the problem, early display adapters used special dual-ported RAM chips - known as video RAM or VRAM - which allow two accesses to take place at the same time. VRAM is more expensive than ordinary dynamic RAM (DRAM). This is a significant factor in the highly competitive video card market, so low and mid-priced display adapters generally use standard memory chips and force the processor to wait if it tries to access a memory location needed by the video controller. This will have an effect on performance, so many top specification video cards still use VRAM despite the extra cost.

The higher the resolution and the higher the refresh rate, the greater the frequency of memory accesses made by the video controller. As we've already seen, at 72 Hz and a resolution of 1024 x 768 the dot clock for each pixel is 70 MHz. A colour depth of 16-bits would require a 16-bit data access each time. DRAM is unable to deliver that sort of access time, and so the video adapter designer must use similar tricks to those used in the PC itself to speed up access to video memory.

The latest video controllers have a 32-bit data path between the controller and video memory, which improves the performance and cuts down the frequency of memory access cycles. Beyond this, the memory on the board can be arranged in banks, and used in page interleave mode. This means that the first 32-bits can be read from one bank, the next 32 from the next bank, and so on. The frequency of accesses to one memory bank is thereby halved, allowing the ceiling on display resolution and refresh rate to be raised.

When specifying a display adapter you need to check whether you need extra memory in order to use the refresh rates that a card claims to support. The manual that accompanies a card may not make this clear. Even the configuration software that comes with a graphics card may not always prevent you from selecting a refresh rate that won't work. If you notice flicker on the display that ought not to be there at the refresh rate you've chosen, insufficient video RAM is probably the reason.

Performance Issues

The memory capacity of a display adapter may put a theoretical limit on the display resolutions that can be used, but a lower, more practical limit can be imposed by the performance. As an example, going from standard VGA 640 x 480 x 16 colours to 1024 x 768 x 256 colours increases by a factor of five the amount of data needed to represent what is shown on the screen. This means that the system has to do five times as much work to get information on to the screen or, to put it another way, performance will be five times worse.

This would not be such a problem if standard VGA graphics mode was acceptably quick to start with but on most machines that is not the case. Fortunately, a number of recent developments are resulting in poor graphics performance becoming a thing of the past.

The standard PC graphics functions are somewhat primitive. Everything that appears on the screen must be laboriously plotted, pixel by pixel. This means that large numbers of bytes must be written by the processor into video memory. In the case of a complete 1024 x 768 x 256 screen - not a particularly high resolution by current standards - 768 KB of data would need to be written to fill the screen.

Unlike the PC's main memory, video memory is not directly connected to the main processor's data bus. Even if the video adapter is built on to the motherboard the memory is still accessed through the expansion bus. In the case of the ISA bus, this limits the rate at which data can be written to the frame buffer to about 2 MB/sec, so to repaint the screen of a 1024 x 768 x 256 graphics display would take a third of a second, even if nothing else was slowing things down. With higher resolutions and greater colour depths it will take even longer. The ISA bus is a bottleneck.

Poor graphics performance was the main reason behind the development of the local bus. By running the bus at the speed of the CPU rather than the 8 MHz of the ISA bus, an increase in

"The latest video controllers have a 32-bit data path between the controller and video memory which improves the performance and cuts down the frequency of memory access cycles."

PC Display Adapters

throughput of from three to six times can be achieved. The 32-bit data path of the local bus offers the potential for a further doubling of throughput.

But although it offers a great opportunity for advertising hype, local bus is not the panacea for all graphics ills. Whatever appears on the screen must still be computed and then written into the frame buffer by the CPU, so graphics performance becomes processor-limited if it is not bus-limited. And the much-vaunted benefits of a 32-bit bus are only realised if you use 32-bit display drivers, which most Windows and DOS drivers aren't. A good ISA graphics card can still outperform a poor local bus one.

Accelerating Graphics

The solution to the problem of poor graphics performance is to build intelligence into the graphics card so that some of the work involved in generating the patterns of pixels that appear on the screen is done by the card and not the system processor. Performance is improved both by cutting down the workload of the CPU and by reducing the volume of data that needs to be transferred over the bus.

There are two types of intelligent graphics card. The oldest type is the graphics co-processor. This is essentially a microprocessor in its own right, though with an instruction set optimised for graphics use, and has its own program memory as well as the video frame buffer.

The programmability of the graphics co-processor means that it can be made to support graphical functions needed by a particular application. However, many co-processed graphics boards are built around co-processors made by Texas Instruments. These chips support a standard set of graphics functions, the Texas Instruments Graphics Architecture (TIGA), and many applications simply make use of these functions.

Graphics co-processors tend to be expensive: they are, after all, complete computers on a board. They also tend not to be either register or BIOS compatible with PC display adapters, and consequently don't support standard PC display modes like VGA. This means you can't even boot into DOS.

To overcome this limitation it's possible to install a graphics co-processor and a VGA card side by side, each with its own monitor. In certain specialist applications this may be a convenient configuration, letting you have one screen for text, another for graphics. However for most people, having two monitors is far from convenient, and so the co-processor card is usually able to accept the output from a standard VGA card, which is routed through to the single monitor when ordinary PC software is in use.

Most co-processor boards obtain the VGA signal via an internal connection to the feature connector on the VGA card. This connector may look like a small version of an ISA expansion slot connector, formed by a cut-out in the board. However, more modern boards use a 26-way pin type connector, the functions of which have been standardised by the Video Equipment Standards Association (VESA). Some cheap VGA cards may not have a feature connector at all, so it's worth checking for one if you plan to install a co-processed graphics card in a machine.

> *"Most co-processor boards obtain the VGA signal via an internal connection to the feature connector on the VGA card. This connector may look like a small version of an ISA expansion slot connector."*

A lower-cost alternative to the graphics co-processor is the fixed function graphics accelerator. An early example of this is the IBM 8514/A adapter. The 8514/A card does not have built-in VGA compatibility either, and so must be used in conjunction with a separate VGA card. Some third-party 8514/A compatible cards do include on-board VGA support, however, and so can be used on their own.

Windows

Today, of course, the most common example of a fixed function graphics accelerator is the Windows accelerator, so called because its accelerated functions are restricted to those most useful under Microsoft Windows. There is a bewildering variety of these to choose from, based on one of a number of graphics accelerator chip sets from several different manufacturers, and offering widely varying levels of performance.

Graphics accelerator capabilities are described in terms of the functions they are designed to accelerate. The ones you'll most commonly find are described in the box "Graphics Accelerator Features Explained".

Software can only take advantage of graphics acceleration if it has installable drivers that provide separate function calls for different graphics operations. If a program clears a window by setting each pixel in the area to white, for example, there is no way the graphics card can recognise that this is an area-fill operation and take over the job. The program has to call an area-fill function, so that the driver software, knowing that the graphics card it is designed to support has this capability, can issue the appropriate commands to it.

Windows provides a whole range of such functions in its Graphics Device Interface (GDI), which is why it is an ideal environment for use with graphics accelerators. However, it is important to realise that you will only get the benefit of a graphics accelerator card when you use the drivers supplied with it. If you use the standard VGA or SuperVGA drivers with a Windows accelerator card, it will treat

it as a standard VGA or SuperVGA card, and you'll get no performance improvement whatever.

Evaluating Performance

It's impossible to judge the performance of a display adapter by simply looking at its specification. Even when comparing Windows accelerators, it is not sufficient to look at the list of graphical operations that are supported by the hardware. The speed with which they are carried out varies from one board to another. The display drivers themselves can also have a big impact on performance.

Under Windows, it is relatively easy to test graphics performance because the Windows GDI consists of a set of graphics functions that can be individually tested.

Performance

Performance in text mode is no longer a real issue with today's 486-based PCs, but graphics performance is another matter. The difficulty in measuring it is that there is no standard interface; performance will depend to some extent on the graphics library used to create each individual application. About the only things that can be measured are pixel read/write operations. A useful program which can be used to test monitor performance, as well as display adapter compatibility and both text and graphics performance is DisplayMate, from Sonera Technologies. The performance of the test PC has a major impact on any graphics benchmark result, and so when comparing the performance of graphics cards it is important that the tests have been carried out under identical conditions, including the refresh rate. Advertised WinMark results are of little value even if the speed of the system used for testing is given.

Few general guidelines can be given, though it is true to say that DOS text and graphics performance is frequently a casualty in the quest for better Windows performance. Notable exceptions to this rule are display cards based on the Tseng Labs ET4000AX/W32 Windows accelerator chip. The ET4000AX/W32 is not only one of the faster fixed function graphics accelerators currently available, but it is fully hardware compatible with the earlier ET4000AX Super VGA video controller. This was one of the best DOS text and graphics performers. However, an even more important benefit if you want a Windows accelerator that will also work well with non-Windows graphics software is that existing ET4000 drivers can be used. This makes ET4000AX/W32 based cards probably the best-supported graphics cards you can currently buy.

Non-Windows Cards

If you are specifying a display adapter for use with non-Windows software then obviously it is important to ensure that drivers for that software are available. It is also a good idea to find out where you can get updated drivers from - not always easy if you have a PC containing a graphics card made by a nameless OEM.

Windows display drivers are difficult to write, and bugs are not uncommon.

There is a problem when using Borland's Paradox for Windows or Quattro Pro for Windows with certain S3 86C801 or 86C805 based display adapters. Shortly after loading the applications, strange effects will be noticed, the cause of which is rapidly dwindling Windows resources. Borland's claim that it a problem with the graphics card drivers has a ring of truth, since the standard Windows drivers and other graphics cards don't experience the problem. On the other hand, only these Borland applications are affected. This particular problem has caused many people hours of wasted time.

Conclusions

When selecting a display adapter, you must take care that it supports the resolutions and refresh rates you want. Make sure that the memory fitted is adequate, particularly if you want to use high colour depths and ergonomic refresh rates. If you think there may be a need to use higher resolutions later, make sure that the card can easily be upgraded with extra RAM.

If you want to use 18-bit or 24-bit colour, do check that performance will be adequate. Many cheaper cards offer at least 18-bit colour, but speed will usually be too sluggish for serious use.

Check the availability of drivers for the software that will be used with the adapter, or you will be forced to use it as a standard VGA card and lose the benefit of any additional capabilities it might have. To make sure there aren't any crippling bugs, test the card and its drivers with the software before making any volume purchase commitments. Find out where updated drivers can be obtained, in case you later discover you need them.

If you're choosing a display card with a view to upgrading existing PCs, test it with the monitors it will be used with to make sure that they can support the new display modes, and that there won't be a need for inconvenient readjustment of the position and size controls every time the user changes display modes.

Given the bewildering choice of display adapters now available, choosing the right one for your users' needs is not easy. But it's worth the effort. The PC no longer deserves its reputation as an inferior platform for graphics.

PCSA

The Author

Julian Moss is a PC writer, journalist and consultant and a regular contributor to PC Support Advisor.

Windows .INI Files

The Windows startup files - WIN.INI and SYSTEM.INI are the key to an efficient Windows operation.

In brief, INI files are ordinary text files which contain customisation options, user preferences and configuration data. WIN.INI mainly contains parameters which affect the appearance and operation of Windows - the user's preferred colour scheme and wallpaper, for example. SYSTEM.INI contains data relating to the hardware and the environment in which Windows is running. The distinction is not a rigid one, however.

In both files, many of the entries are tied to options which you set in the Control Panel. If you use Control Panel to install or remove fonts, for instance, your choices will be recorded in WIN.INI. In most cases, the change comes into force as soon as you finish with the relevant dialogue. If it is necessary to restart Windows in order for the change to take effect, Control Panel will offer to do this for you.

Other parts of Windows also update the two files. Windows Setup will alter SYSTEM.INI to reflect changes in installed drivers, while File Manager uses WIN.INI to record any new file associations. As well, many applications maintain their own sections within WIN.INI, typically to record user preferences that only affect the application.

Since both files are ordinary text files, it is also possible to edit them directly, using an ordinary ASCII editor. Indeed, many of the options within the files can only be altered in this way. When you alter the INI files manually, you must always restart Windows for the changes to take effect.

File Format

The entries in WIN.INI and SYSTEM.INI follow a consistent format. Related entries are grouped into sections, each of which begins with a header enclosed in square brackets. Thus, [desktop], [sound] and [Windows Help] are all valid section headers.

Within a section, each item occupies a separate line. These lines have the general format:

```
keyword=value
```

The following entry, for example, sets the size of the icon titles to 10 points:

```
IconTitleSize=10
```

The value can be a number, a string, or a boolean value. Boolean values are usually 0 (false) or 1 (true), but you can also type Yes, No, True or False.

Here is a longer example, taken from a typical WIN.INI file:

```
[windows]
BorderWidth=3
KeyboardSpeed=12
CoolSwitch=False
Programs=com exe bat pif
device=HP LaserJet III,hppcl5a,LPT1:
```

The order of the sections is not significant, nor is the order of lines within a section. Case is also irrelevant. You may add blank lines to improve readability.

If a line begins with a semi-colon, it is a comment. Rather than deleting an unwanted line, it's a good idea to comment it out to make it easier to re-instate if the need arises.

Default Values

In most cases, if the value after the equals sign is omitted, a default is used. For instance,

```
IconSpacing=
```

sets the icon spacing to 77, since this is the default. The same default applies if the entire line is missing. There are exceptions, though. If certain lines are missing from SYSTEM.INI, Windows will refuse to load.

Editing The Files

Although you can edit INI files with any ASCII editor, you might prefer to use the special System Editor which comes with Windows (Figure 1). To do so, run SYSEDIT.EXE from Program Manager's File Run dialogue. The editor automatically opens WIN.INI and SYSTEM.INI - and also AUTOEXEC.BAT and CONFIFG.SYS - each in its own window. When you save a file, a backup is created (with the extension SYD). There are a number of more specialised INI files browsers and editors on the market. One is INI Navigator, a shareware program by Robert Di Bacco (Figure 2). It lets you browse and edit any INI file in the Windows directory, using an index of sections as an aid to navigation. The program optionally creates backup files, which you can browse and selectively restore. *[This program will be on PCSA Utility Disk 64 - Ed.]*

Recovering A Lost WIN.INI

If you accidently delete WIN.INI - or some vital part of it - you won't have to re-install Windows in order to get it back. If part of the file still exists, save it under a different name. Then expand the original "source" file, which you will find on Disk 1 (Disk 3 if you have Windows for Workgroups). Insert the disk in Drive A:, and type the following from DOS:

```
expand a:win.sr_ c:\windows\-
win.ini
```

This will create a bare-bones WIN.INI file. Next, open the file in a text editor (outside Windows), and delete the "SetupWin=1" line from the [Windows] section. If this line was not deleted, Windows would think that it was in the middle of the installation routine, which is not what you want.

Save the file, exit the editor and start Windows. You might notice some unusual fonts in the Program Manager window, but this is easy to correct. Go to the Printers section of Control Panel and re-install your printers. Then go to the Fonts section, select the System directory, and press the Select All button to re-install the default fonts. Exit Control Panel.

If some of the original WIN.INI file is salvageable, you can copy any valid sections of it to the new version. You might not be able to recreate the entire file in this way, but you will get enough of it back to get you started again.

Cleaning Up

One problem with INI files is that they are used as something of a dumping ground by Windows applications. When you install a new word processor or spreadsheet, it might write a lot of information in one or both files - usually WIN.INI. If you later remove the application, the entries remain. These unwanted entries do no harm, but a bloated INI file causes Windows to load more slowly and to consume more resources.

Often you will be able to recognise the unwanted text so that you can delete it - but not always. One solution is to make a copy of the INI files before installing a new application, then use the DOS file-comparison command to identify the new entries. For example:

```
fc win.ini win.sav >changes.txt
```

This would record the differences in a file named CHANGES.TXT. I've also seen people using their word processor's "redline" feature to document the differences.

A more elegant solution is provided by Rainer Baltruweit's WinClean, a shareware program. *[This will be on PCSA Utility Disk 64 - Ed.]* You start it before you install the new application. It keeps a record of any directories which the install routine creates, any new files in the Windows and System directories, and any new entries in WIN.INI. When you no longer want the application, it will delete all these items en masse.

The Files In Detail

Let's now look at the contents of the INI files in detail, starting with WIN.INI. I'll describe the various entries in functional groups, rather than on a section-by-section basis (see also the text box, "WIN.INI Sections"). There are over 300 entries that can appear in the INI files so there's no room in this article to concentrate on any but the more important ones.

List Of Ports

The [ports] section of WIN.INI establishes a list of the printer and communications ports recognised by Windows, for example:

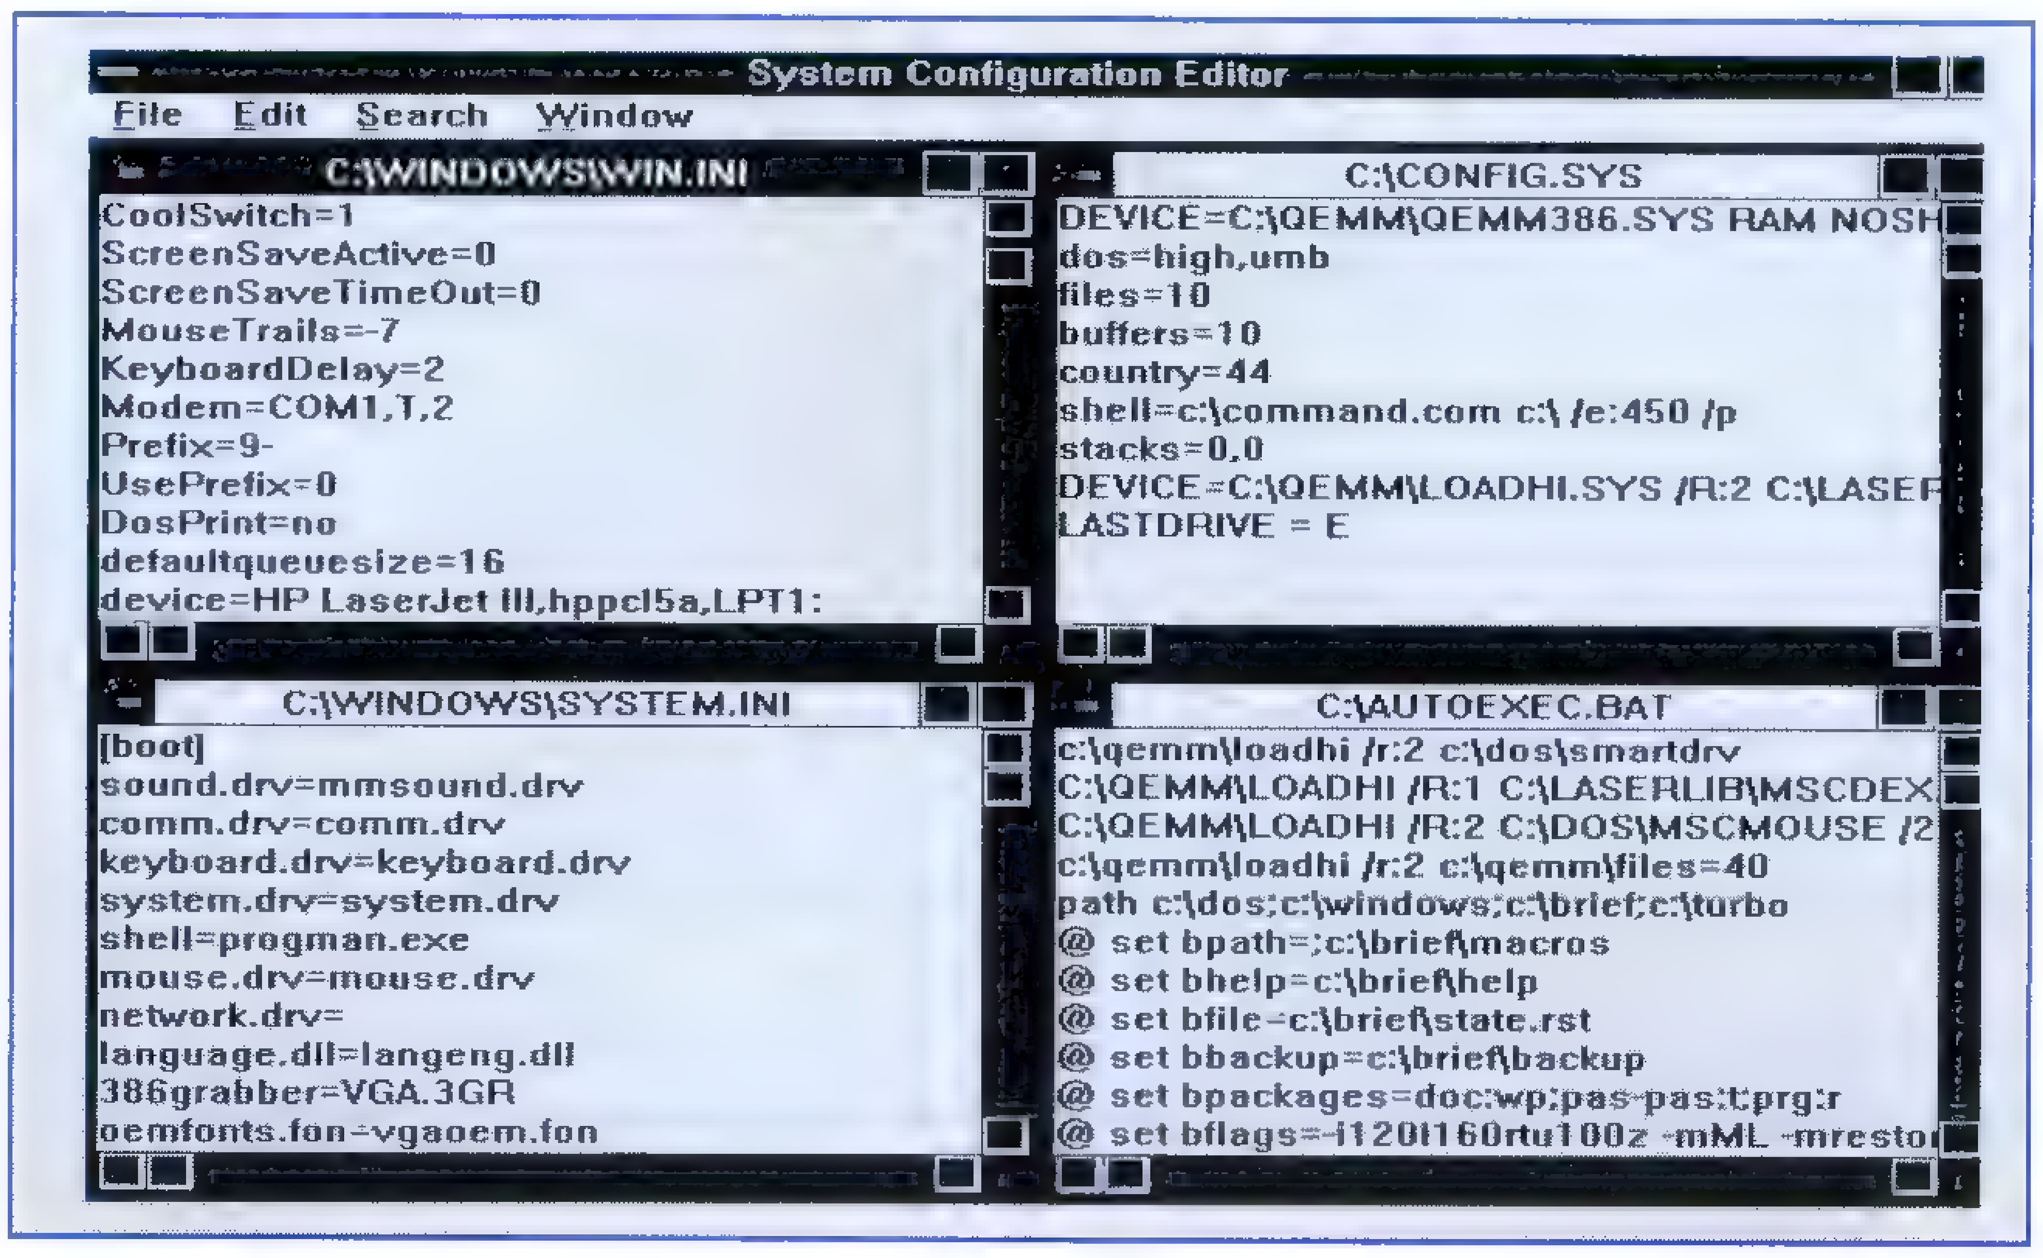

Figure 1 - Use the Windows System Editor to edit the "big four" startup files

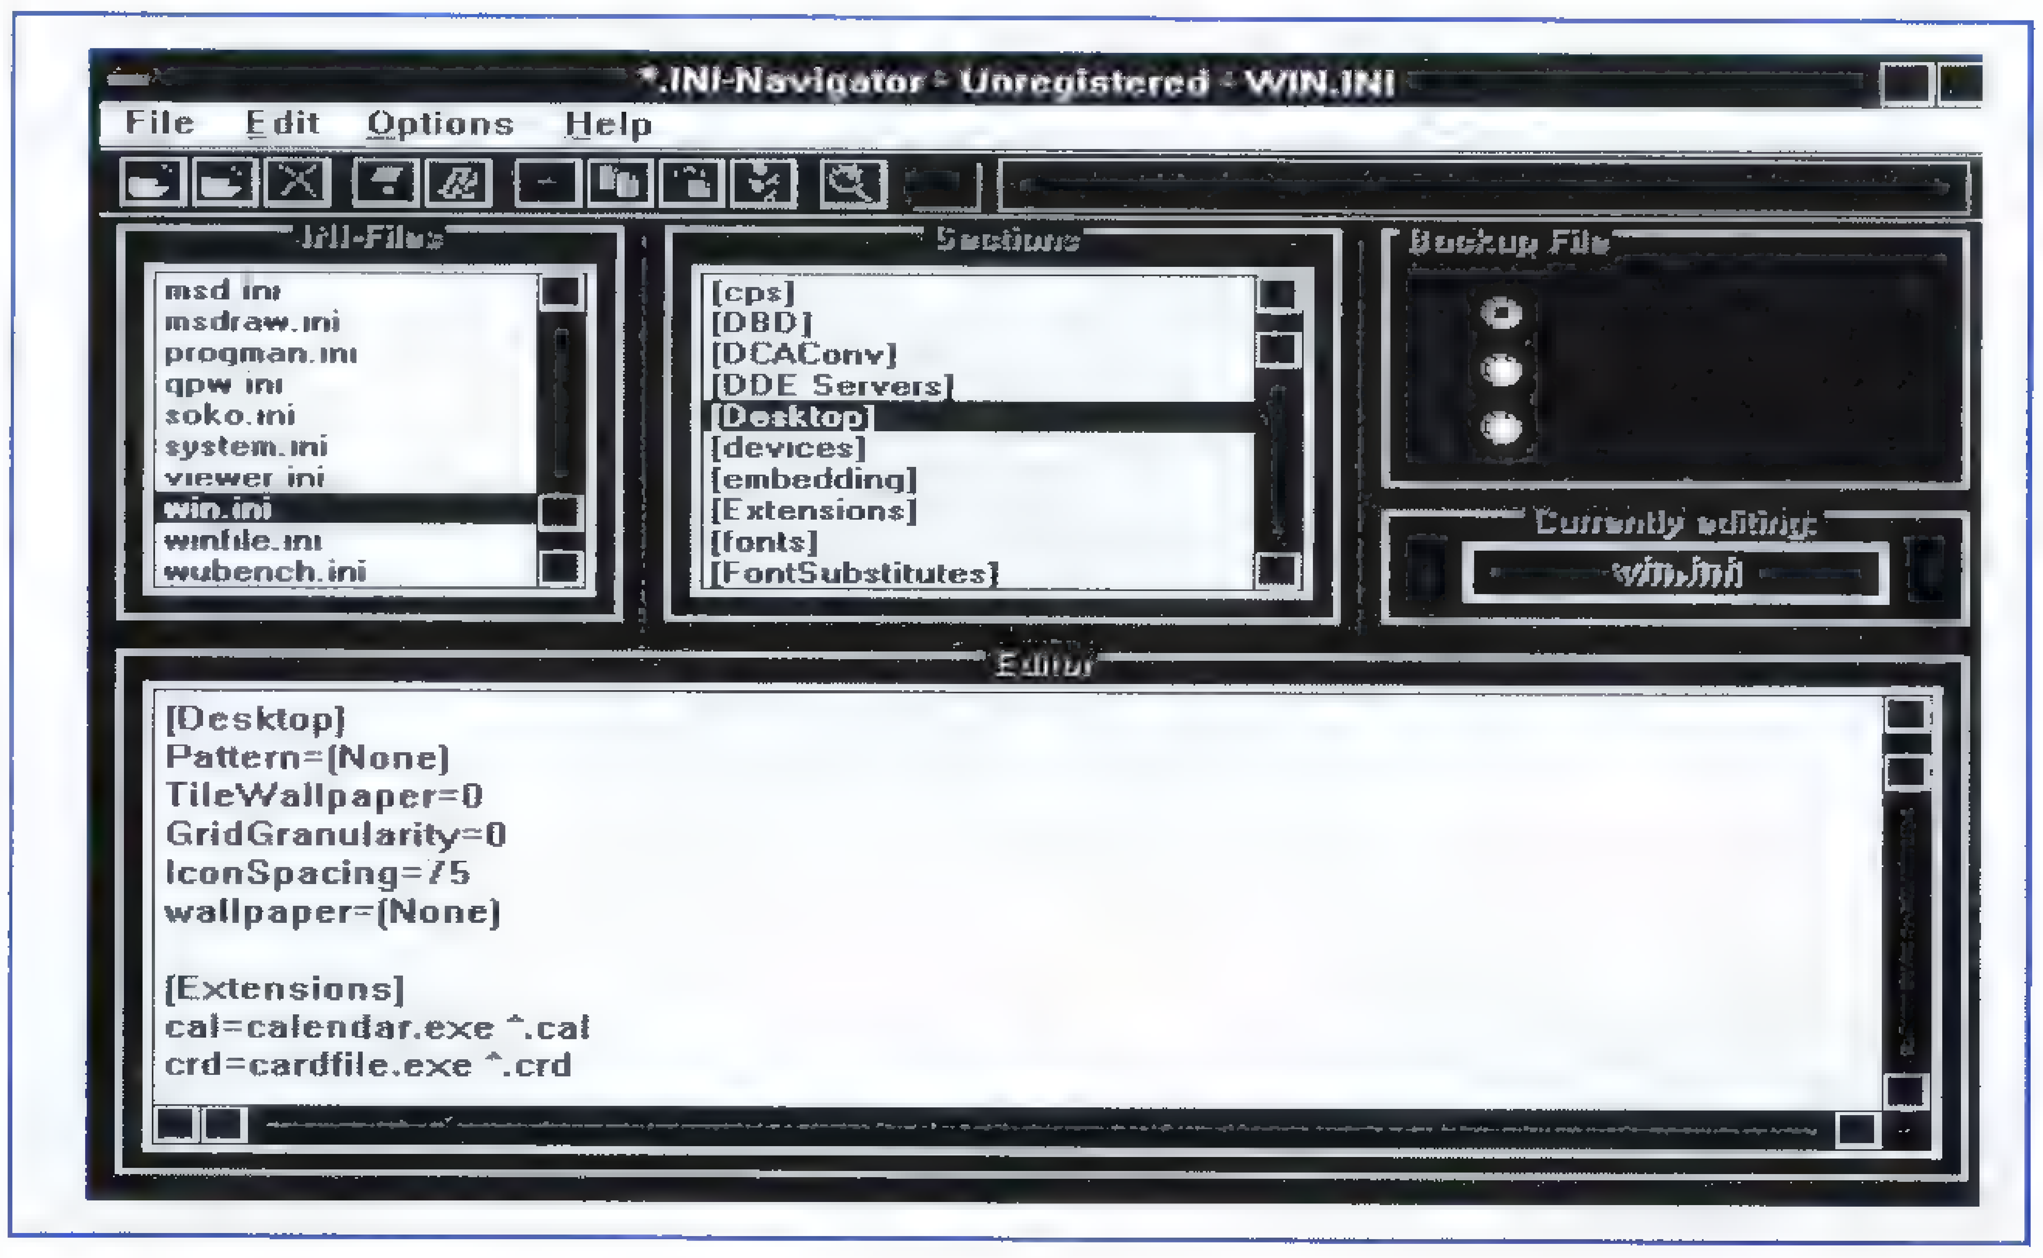

Figure 2 - The INI Navigator is one of many specialised INI editors available

Windows .INI Files

LPT1:=
LPT2:=
COM1:=9600,n,8,1
COM2:=9600,n,8,1
EPT:=
FILE:=
LPT1.DOS=

The serial ports are followed by their communication parameters, in the same format as the DOS MODE command. You would normally use Control Panel's Ports dialogue to alter these parameters. The EPT entry refers to IBM's Enhanced Printer Port.

The FILE entry is used to direct printed output to a file. If you assign a printer to this "port", Windows will prompt you for a filename each time you print. You can also use a specific filename here, for example:

c:\work\myfile.prn=

This directs printing to the specified file (note that there is no colon).

The final entry in the above example is used to bypass Windows's normal parallel-port output. Because there is no colon, Windows thinks that this refers to a file, so it sends the output to DOS. DOS sees the first four characters of the port name as a device (eg, LPT1), and prints accordingly.

Assigning Printers To Ports

Having established the list of ports, you can now turn to the [PrinterPorts] section to see which printers are assigned to them. This section contains entries in the following format:

printer=driver, port, device-timeout, retry-timeout

For example:

HP LaserJet III=hppcl5a,LPT1:,15,45

You would not normally edit this section. You use the Printers module in Control Panel to assign printers to ports and to define their time-out parameters. To find the current default printer, look for the following entry in the [Windows] section:

Device=printer-name, driver, port

For example:

Device=HP LaserJet III,hppcl5a,LPT1:

Still in the [Windows] section, the following entries establish the default values for the two timeout parameters in the [PrinterPorts] entries:

DeviceNotSelectedTimeout=number_of_seconds
TransmissionRetryTimeout=number_of_seconds

Fonts

The currently installed fonts are listed in the [fonts] section. Each entry is in the form:

font-name=font-file

The font name is the name of the font as it appears in menus. For TrueType fonts, the font file is an information file (extension FOT) which contains a pointer to the actual font file (extension TTF). For other fonts, the font file is the font itself (extension FON). Some examples:

Arial (TrueType)=ARIAL.FOT
Courier 10,12,15 (VGA res)=COURE.FON
MS LineDraw (All res)=WINLD.FON

You can edit this section yourself, but adding an entry does not make the font available. To do that, you must install it via the Fonts dialogue in Control Panel. The [fonts] section does not list printer-resident fonts. The names of these fonts generally appear in a separate file, the details of which vary according to the printer driver.

The [FontSubstitutes] section is used to equate installed fonts with fonts which are not currently available. For example, if you have a document formatted in Univers, and if Univers is not installed on your system, you might add the following entry to display the font as Arial:

Univers=Arial

Don't confuse the [FontSubstitutes] section with the PostScript driver's font substitution table. If you have edited the substitution table in the PostScript Advanced Options dialogue (in Control Panel), your WIN.INI file will contain a [PSCRIPT] section which contains entries like the following:

Futura=Helvetica

This tells the driver to use the printer's Helvetica font for text which has been formatted with Futura. A value of 0 after the equals sign means that the font should be downloaded as a soft font.

TrueType Settings

The [TrueType] section contains four entries which govern the way that Windows uses TrueType.

TTEnable=0 or 1
TTOnly=0 or 1

These entries correspond to the "Enable TrueType fonts" and "Show only TrueType fonts in applications" checkboxes in the TrueType dialogue in Control Panel's Fonts module. The defaults are 1 and 0 respectively

OutlineThreshold=number

This entry determines the method which Windows users to render fonts. If you have trouble printing large or complex fonts, try decreasing the value of the number to around 100 (default: 256). To do so, edit the file directly.

TTIfCollisions=0 or 1

If you have a TrueType and a non-TrueType font with the same name, Windows will normally use the non-TrueType font. If that is not what you want, change the value in this entry from 0 (the default) to 1. Edit the file directly.

Network Settings

The [networks] section contains the following entries:

drive=network-server-and-share

This entry specifies which network connections to restore at the start of

each session. You can alter it via the Network Connections dialogue in File Manager's Disk menu.

InRestoreNetConnect=0 or 1

If this entry is set to 1 (the default), Windows restores the previous network connection at startup. It corresponds to the Reconnect option in Control Panel's Networks dialogue. If this dialogue does not have a Reconnect option for your network, edit the file directly.

port=network-printer-path

This entry defines the path and port for a network printer. You can alter it by using the Network dialogue within the Printer Connect dialogue in Control Panel.

By default, Windows will display a warning message on start-up if it is configured for a network and the network is not available. If you don't want this warning, locate the following entry in the [Windows] section, and change the value to 0:

Netwarn=0 or 1

You can also alter this setting via Control Panel's Networks dialogue.

Keyboard And Mouse

The [Windows] section contains a number of commands which govern the operation of the keyboard and mouse.

KeyboardDelay=milliseconds
KeyboardSpeed=milliseconds

These entries, which correspond to the two sliders in Control Panel's Keyboard dialogue, govern the keyboard repeat rate. The first specifies the delay before a key starts to repeat (default: 2); the second is the interval between repetitions (default: 31).

DoubleClickSpeed=milliseconds

This entry specifies the maximum time between the two clicks in a double-click (default: 452). If you have difficulty in double clicking, try increasing the value. You can alter it via the Mouse dialogue in Control Panel.

DoubleClickHeight=pixels
DoubleClickWidth=pixels

These entries specify the amount by which the mouse pointer is allowed to move between the two clicks in a double-click (default: 4 in each direction). If it moves by more than the stated number of pixels, Windows regards the action as two single clicks. To alter the values, edit the file directly.

MouseTrails=number

You use the Mouse Trails checkbox in the Mouse dialogue to switch mouse trails on and off, but you must edit this entry directly if you want to vary the number of pointers which form the trails. Permitted values are between 0 (no trails) and 7 (the default value when mouse trails is switched on).

SwapMouseButton=0 or 1

Set this to 1 to swap the left and right mouse buttons. You can alter it from the Mouse dialogue.

MouseSpeed=0 or 1 or 2
MouseThreshold1=pixels
MouseThreshold2=pixels

These three entries together determine how the mouse acceleration works. If MouseSpeed is 0, the mouse pointer moves at a constant speed relative to the mouse. If it is 1 (the default), the pointer doubles in speed in relation to the mouse each time that it moves more than MouseThreshold1 pixels (default: 5) in one movement. If MouseSpeed is 2, the pointer quadruples in speed in relation to the mouse each time that it moves more than MouseThreshold2 pixels (default: 10) in one movement. Adjusting the Mouse Tracking Speed slider in the Mouse dialogue will provide a simpler way of altering the acceleration.

WIN.INI Sections

These are the WIN.INI sections which you are most likely to come across.

[colors]	Corresponds to Control Panel's Color dialogue. Note the US spelling.
[compatibility]	Deals with minor incompatibilities between Windows 3.0 and 3.1. Do not edit or delete it.
[desktop]	Governs the appearance of Program Manager and the background screen.
[devices]	Lists active output devices, for compatibility with Windows 2.x. Do not edit or delete it.
[embedding]	Contains a list of OLE objects. It duplicates the data held in the registration database, and is included for compatibility with Windows 3.0. Do not edit or delete it.
[extensions]	Associates file extensions with applications.
[fonts]	Maps font names to font files.
[fontSubstitutes]	Equates installed screen fonts with fonts which are not available.
[Intl]	Corresponds to the settings in Control Panel's International dialogue.
[MCI extensions]	Associates file extensions with Media Control Interface drivers. It should never require editing.
[network]	Contains network-related settings.
[ports]	Lists the available printer and communication ports.
[PrinterPorts]	Lists installed printers and their corresponding ports.
[programs]	Lists paths to programs which are required when documents are launched.
[PSCRIPT]	Contains the PostScript driver's substitution table.
[sound]	Corresponds to the settings in Control Panel's Sound dialogue.
[TrueType]	Contains TrueType-related settings.
[Windows]	Contains basic settings that affect the environment and the desktop.
[Windows help]	Contains customisation options for the Help applet.

Windows INI Files

Extensions & Associations

The [extensions] section stores the file associations that you make with the Associate command within File Manager. For example, if you have associated the XLS extension with Microsoft Excel, the [extensions] section might contain:

xls=c:\excel\excel.exe ^.xls

Then, when you "launch" an XLS document from File Manager or Program Manager, Windows will load Excel which will in turn load the document (the caret in the above entry is a place-holder for the filename).

Although you can add entries to this section yourself, you rarely need to do as most applications create their own associations. But you will probably want to prune this section to remove entries belonging to applications that are no longer installed.

If, when you launch a document, Windows cannot find the application, it prompts you for the application's path. It then records the path in the [programs] section of WIN.INI. For example, if you moved Paintbrush to c:\graphics and then tried to launch a BMP file, Windows would ask you where Paintbrush was located. After you had told it, the [programs] section would contain this line:

pbrush.exe=c:\graphics\pbrush.exe

You should never need to edit this section.

You can also define documents by means of the Documents= entry in the [Windows] section, for example:

Documents=PRG PAS C

Files with these extensions will be regarded as documents by File Manager, but cannot be launched as they are not associated with applications. The only reason I can think of for using this entry is if you were filtering a File Manager view to show documents only. To do this, edit WIN.INI directly. The Programs= entry in the [Windows] section lists the file extensions that Windows regards as belonging to executable programs. Its default value is:

Programs=com exe bat pif

You might want to add an extension to the list, for example, to run a screen saver (extension SCR, but it's actually a standard .EXE file internally) directly from Program Manager. Edit WIN.INI directly.

Startup Programs

These two entries list the applications and documents that you want Windows to launch automatically on start-up:

Run=list-of-programs
Load=list-of-programs

The first runs the programs normally; the second runs them as icons. You would normally use Program Manager's Startup group to create these entries. You would create or edit them yourself if you were using another shell, such as File Manager or Norton Desktop for Windows, both of which know about these entries.

Aesthetics

Many WIN.INI entries are provided purely for aesthetic purposes, and have little functional value. Still, users do like to individualise their environments, and these settings provide plenty of scope for doing so.

The [desktop] section governs the appearance of Program Manager and of the background screen. Except where otherwise stated, these items can be altered from within Control Panel's Desktop module.

IconSpacing=pixels

This entry determines the horizontal spacing that separates icons when you choose the Arrange Icons command in Program Manager (default: 77). You can always move icons closer together or further apart by dragging them with the mouse.

IconVerticalSpacing=pixels

This entry is similar to the previous one, except that it governs the vertical spacing. The default depends on the display resolution. To alter it, edit the file directly.

IconTitleFaceName=fontname
IconTitleSize=size

These entries determine the icon titles' typeface and point size respectively (default: MS Sans Serif 8 pt). To alter them, edit the file directly. You can set the typeface to any font with a FON extension which is listed in the [fonts] section (in fact, you can choose any font, but Windows will substitute the closest matching FON file). Increasing the point size is useful if you have trouble reading the titles on a high-resolution display.

GridGranularity=size

This entry establishes an imaginary grid to which application windows are aligned. The grid size is expressed in units of eight pixels (maximum: 49). A value of 0 (the default) turns off the grid. The setting does not affect the placing of icons or child windows.

The [Windows] section also contains entries related to the settings in the Desktop dialogue.

BorderWidth=pixels

This entry determines the thickness of the borders of resizable windows. Permitted values are 1 to 49 (default: 3). Increase the value if you find it hard to select borders on a high-resolution display.

CursorBlinkRate=milliseconds

Increase the value in this entry if you have trouble seeing the text cursor, for example on an LCD screen. Select a lower value if you find the blinking cursor distracting. Within Control Panel, you can choose a value between 200 and 1,200 (default: 550). Edit the file directly to enter values outside this range.

ScreenSaverActive=0 or 1
ScreenSaverTimeOut=seconds

The first of these entries enables or disables the screen saver (default: 0, or disabled). The second entry specifies

the period after which the screen saver springs into life. The name of the screen saver program is specified in the Scrnsave.Exe= entry in the [boot] section of SYSTEM.INI.

CoolSwitch=0 or 1

Although not obvious, this delightfully named setting corresponds to the "Fast Alt-Tab switching" checkbox in the Desktop dialogue. It is enabled by default.

MenuDropAlignment=0 or 1

This one is for users who really want to be different from their colleagues. If you change the value to 1, drop-down menus will be right-aligned with their titles, rather than left-aligned as normal. Edit the file directly.

MenuShowDelay=milliseconds

This entry specifies how long you must wait for a menu to open after you click on a menu item. On 286-based systems, a delay of 400 milliseconds is used to prevent jerkiness when you sweep the mouse over open menus. Reduce the figure (edit the file directly) if you find the delay annoying. On 386s and above, the default is 0.

Desktop Colours

The [Colors] section of WIN.INI corresponds to Control Panel's Color dialogue. Its entries are in the following format:

component=red-value green-value blue-value

The components are the screen elements whose colours can be customised by the user. They correspond to the 21 items listed in the Color Palette section of the dialogue. The three values are integers, in the range 0 to 255, and represent the intensity of the respective colours. For example:

ActiveTitle=0 0 128

specifies that a window's active title appears as mid-blue.

The SYSTEM.INI File

Many of entries in SYSTEM.INI are created when you install Windows or when you add new drivers, and should never be altered directly. I'll concentrate here on the few entries that you might want to alter yourself. In each case, you must directly edit the file to alter these settings.

The following entries are in the [boot] section:

shell=program-filename
Taskman.Exe=program-filename

These entries specify the executable programs which are used as the Windows shell and the task manager respectively. The defaults are PROGMAN.EXE and TASKMAN.EXE. You might want to change the shell to File Manager (FILEMAN.EXE) or to a third-party shell. The task manager is the program list which pops up when you press Ctrl-Esc. Third-party task managers are rare. If you don't need the services of a task manager, you could install any small utility that you want to get at especially quickly.

fonts.fon=font-filename
fixedfon.fon=font-filename
oemfonts.fon=font-filename

These entries govern the fonts used for, respectively, the main Windows proportional system font, the fixed system font used by Windows 2.x applications, and the OEM font used by the clipboard viewer. In each case, the font file has the extension FON. The defaults depend on the display driver.

The following entries are in the [NonWindowsApps] section:

CommandEnvSize=new-size-in-bytes

Each DOS session within Windows inherits the environment space which was available before Windows started (and which is typically established by the shell command in CONFIG.SYS). Use the above entry if you want more environment space under Windows.

MouseInDosBox=0 or 1

If you do not want DOS applications to recognise the mouse, change this entry from 1 (the default if a DOS mouse driver is installed) to 0.

ScreenLines=number-of-lines

This entry specifies the number of screen lines initially displayed in a DOS session (default: 25). The DOS application must be able to support the specified number of lines.

SwapDisk=drive:directory

This entry specifies the drive and directory to which Windows swaps DOS applications in Standard mode. The default is the location specified by the TEMP environment variable.

One final entry worth knowing about is the following, which is in the [386Enh] section:

DosPromptExitInstruc=0 or 1

Change the value in this entry from 1 (the default) to 0 if you don't want to see the large, boxed message that appears at the top of the screen whenever Windows starts a DOS prompt session.

PCSA

The Author

Mike Lewis is a PC consultant and writer specialising in Windows and databases.

How To Choose A CD-ROM Drive

There are dozens of CD-ROM drives on the market, but not all of them will be suitable for you. Wendy M Grossman explains why.

CD-ROM is finally taking off. In 1992 and 1993, the number of discs on the market doubled. At the same time, new hardware has been introduced: writable CD-ROM (CD-R), Photo CD, and even Philips' Compact Disc Interactive (CD-I).

The most obvious initial attraction with CD-ROM is the amount of data that can be held in a small space: 580 MB on a single disc. The newer multimedia applications, however, show off just what can be done with this, mixing text, pictures, sounds and even video clips. Most such titles are, however, aimed at the home market: encyclopedias, introductions to music and musical instruments, and other educational and mass-market titles. The discs currently marketed for corporate use tend to be simpler: reference discs of technical support files, high-end databases, dictionaries, clip art collections, or commercially distributed software.

Disk Types

What kind of discs a user needs to run is important in both specifying and supporting a CD-ROM drive - and the machine attached to it. The crucial point is the type of search software the disc uses and the type of data it holds, since the drive itself can be hooked to any old PC that has a slot to take the necessary interface card. Such a system is adequate for character-based discs since these run under DOS.

However, even with these relatively simple discs, it's important to bear in mind that many discs require that their search software be installed on the PC's hard disk rather than running from the CD-ROM. Even some discs that don't require that you set up a special directory to hold temporary files. This all adds to the support burden.

Newer discs like the Complete Oxford English Dictionary on CD-ROM have Windows-based search software that demands 4 MB of RAM and 2 MB of hard disk space - more if users need to store searches. The newest multimedia discs demand better-specified drives, 5 MB or more of hard disk space, sound board, and runtime video drivers.

If users are going to run such discs, the best guideline for PC specifications is the Multimedia II specification recommended by Microsoft:

486SX/25 MHz processor
4 MB of RAM
160 MB hard drive
1.44 MB 3.5in floppy drive

"A few manufacturers still sell drives that adhere to their own proprietary formats. Philips' CD-I is one example; also proprietary are a couple of Sony designs that are aimed at the consumer market."

Also, of course, a CD-ROM drive that's XA-compliant, multi-session-capable, and double-speed (300 KB per second transfer rate). Some discs actually scan the PC system and complain if the PC isn't up to the recommended specification. Also check video requirements: some newer discs demand 256-colour SVGA resolution.

Photo CD

Photo CD is even more resource-hungry, especially if you want to manipulate large images in 24-bit colour. Kodak recommends the following minimum specification for a PC to run Photo CD and image manipulation software:

486 processor
16 MB RAM
25 MB of hard disk space
Windows

You can get away with less, but you'll find working that way very slow.

Standards

Complicating matters is the introduction of new standards. Where a couple of years ago almost all discs and drives adhered to the High Sierra, also known as ISO9660, standard for file and directory structure, new developments mean that many new discs can't be read by older machines. New machines are, however, backwards compatible.

To read High Sierra discs, the attached PC must be running MSCDEX.EXE, Microsoft's CD-ROM driver. This demands some memory overhead, about 35 KB. This driver is regularly updated - the current ver-

sion shipping with DOS 6.2 is 2.23 - and new versions can be downloaded from online services such as CompuServe's CD-ROM forum or come supplied with upgrades to DOS or Windows. Along those lines, it's worth checking a hard disk that has had such upgrades for multiple copies of the file so the ones not in use can be deleted to free up disk space.

Note that MSCDEX tends to be DOS version-specific and it is highly possible that the version supplied with a newly-purchased drive won't work with your particular PCs. Sometimes you can tweak this with SETVER, but it's safer to go get a version of MSCDEX that will work correctly. Corel supplies a replacement for MSCDEX.EXE with its SCSI adaptor, which comes bundled with some CD-ROM drives. This has the advantage that it can be unloaded from memory if necessary.

All CD-ROM drives advertise that they can play audio CDs, which have their own standard, known as the Red Book, or CD-DA (for compact disc-digital audio), standard. Drives vary widely, however, in how convenient it is to do this: most require special software to be loaded on the attached PC, and this software varies widely in flexibility and ease of use. A number of inexpensive shareware utilities exist for playing audio CDs on a CD-ROM drive, and these are readily available online. (If staff are in the habit of playing audio CDs, ensure that speakers and headphones are not placed near the PC, monitor or floppy disks, as they contain magnets.)

"The one disadvantage of using the parallel interface - either using a Trantor adapter or directly - is that it slows down throughput immensely, even if the PC has an enhanced bidirectional printer port."

CD-ROM XA

CD-ROM XA, for extended architecture, is the standard that lays down rules for interleaving audio and visual (pictures or video) data. This interleaving is what makes it possible to synchronise video clips with audio tracks, something that wasn't possible under the previous Yellow Book standard, which specified different modes for audio and visual data. Many new drives are XA-compliant, even though few discs so far require it.

The exception that does require XA compliance is Photo CD, which is coming into its own in the field of commercial graphics after a false start in the domestic market. The idea behind Photo CD, which was jointly developed by Kodak and Philips and released in 1992, is that ordinary photographs can be transferred onto a disc and stored in compressed, digital form. The photographs can then be loaded onto a PC and manipulated, reproduced, or imported into other applications. Because each disc can hold approximately 100 photos, an essential part of the system is that data can be recorded in multiple sessions. This requires a different directory structure from the more common single-session discs, so Photo CDs are not readable by ordinary High Sierra format drives.

Most - but not all - XA compliant drives are multi-session drives. It's important to check this point if the drive is going to be used for Photo CD applications - even commercial collections of photographs create difficulties on single-session drives.

Aside from these widespread standards, a few manufacturers still sell drives that adhere to their own proprietary formats. Philips' CD-I is one example; also proprietary are a

Networking

CD-ROM drives, like most other devices, can be networked, and there are advantages to doing so, if only to streamline installing software - a system administrator should be able to start the installation running and leave it largely unattended, since there's no need to swap discs.

How easy it is to network a drive will depend on your network operating system's built-in facilities. In general, it's simpler with peer-to-peer networks, where the drive simply appears as an extra drive letter on the user's desktop. In other situations, special drivers are needed. The industry is responding with new products intended to make sharing CD-ROM drives easier, for example, Microtest Discport, designed for NetWare systems, or LANshark Systems' CD-Direct, for Banyan Vines networks.

Besides the technical installation issues, it's important to consider how to handle user access to discs. One solution where everyone needs access to the same small set of discs is a drive that holds a magazine - Pioneer, for example, makes a six-disc "jukebox". Each of the six magazine holders gets assigned its own drive letter, so it appears to the user as if there are six extra drives. This set-up is easy for users and works well, but it's important to ensure that the discs stay in the same place in the magazine - otherwise the search software has to be reconfigured to find the disc on its new "drive". With some discs, this can mean reinstalling the search software. One way to handle this problem is to have extra magazines, and install the discs in them permanently, switching whole magazines rather than individual discs.

CD-ROM

> *"Newer discs like the Complete Oxford English Dictionary on CD-ROM have Windows-based search software that demands 4 MB of RAM and 2 MB of hard disk space - more if users need to store searches."*

couple of Sony designs that are aimed at the consumer market. The discs for these machines cannot be read by standard machines, or vice-versa.

CD-R

Kodak is investing heavily in Photo CD technology, building on it with image database software, formats for Photo CD catalogues, medical applications, and the blending of photographs, text, graphics, and sound. Its most recent development is CD-R, which can record CD-ROM discs that can be read in any standard drive. These machines are still rather expensive - about $5,000 or so - and you may need to add in the cost of formatting or authoring software. Blank discs, however, are inexpensive, working out to about $0.05 per megabyte, which is actually cheaper than less durable media like tape or floppy disk. It's expected, therefore, that as CD-R drives get cheaper, this will become a popular way to archive and distribute data internally.

1993 saw significant speed improvements. Standard-speed drives run at the same speed as audio CD players, and transfer data at approximately 153 KB per second. To meet the demands of graphics, clearly more speed was needed. There aren't too many ways of increasing the speed: change the size of the pits in the disc's surface, which would require expensive manufacturing changes, or increase the physical speed of the drive so the disc spins faster. The latter is the way the industry is going, and a number of companies, including Philips, NEC, Hitachi, and Toshiba, released double-speed drives in 1993. Pioneer has gone one better, and released the first quad-speed drive. A number of manufacturers are planning to launch triple-speed drives in 1994.

Speed

In general, the faster the drive the better. It isn't critical, though, for the text-based discs. Where it starts to matter is on discs with video clips; if the drive is too slow, the movement becomes jerky, halting while the disc reads the next lot of data. The synchronization of audio and video can also be affected.

Tests generally show little difference between standard-speed drives. This is less true with double-speed drives, where on-board caching can make significant differences in speed, as much as 10% in some cases. But bear in mind that the fastest drive on the market - Pioneer's quad speed - is still at least five times slower than an average hard drive.

One way of speeding up disc access times is caching software, and products are coming on the market to handle CD-ROM drives, including the latest version of Symantec's Norton SpeedCache Plus, and the new version of SMARTDRV.SYS bundled in DOS 6.2.

Interface

Speed is also affected by the type of interface that hooks the CD-ROM drive to the PC. There are three choices: SCSI, parallel, and proprie-

Care For CD-ROM Discs

CD-ROM discs are much more vulnerable to damage than audio CDs - audio CD players can "repair" small amounts of damage to discs on the fly by averaging the bits of information surrounding the errors. This isn't possible with computer data, where the integrity of each bit is crucial. Besides, audio CDs are much cheaper to replace.

Accordingly, it's important to protect your investment in CD-ROM discs by treating them carefully. The first and most important enemy is dust. Caddy-based CD-ROM systems are a nuisance compared to the convenience of front-loading direct systems, but the caddies protect the discs, especially if you buy spare caddies for the most often-used discs and use them as permanent protectors. Such a set-up also minimises handling and, therefore, fingerprints.

In addition, some of the better drives come with built-in dust protection, such as double doors and automatic lens cleaners. These are worth looking for in particularly dusty environments. Also important in some situations is the ability to lock the disc into the machine: some drives come with DIP switches that can be set to disable their eject buttons.

None of these protections can be used with the portable CD-ROM drives, which are the very ones that most need the protection. These drives are generally either top or front loading, and don't use caddies. They are, though, very convenient, since they connect to the PC via the parallel port. They are, of course, relatively slow.

> *"Corel supplies a replacement for MSCDEX.EXE with its SCSI adaptor, which comes bundled with some CD-ROM drives. This has the advantage that it can be unloaded from memory if necessary."*

tary. Many drives still come with proprietary 8-bit or 16-bit cards which fit into a PC's slot and control the drive. These tend to be cheaper, but choosing a drive with a SCSI interface allows more flexibility.

Besides the obvious advantage that you can chain other devices to the SCSI adapter, using a SCSI interface also opens up the possibility of connecting the CD-ROM drive to a notebook via a SCSI-to-parallel adapter. These are made by Trantor, which also makes SCSI adapter cards, and are available through CD-ROM specialist dealers.

The one disadvantage of using the parallel interface - either using a Trantor adapter or directly - is that it slows down throughput immensely, even if the PC has an enhanced bidirectional printer port. But so far there are no third-party CD-ROM drives designed to work with the proprietary slots that laptops and notebooks have, although there are a few portables with CD-ROM options. Some docking stations, however, come with SCSI ports.

Installation

One area where manufacturers of both drives and discs tend to do poorly is installation instructions and documentation. No one has yet developed an intelligent installation routine which scans the system and determines which interrupts, DMA channels, and memory addresses are free. The default settings each manufacturer provides generally work on PCs that are free of other devices, but PCs already loaded up with sound boards, scanners, network cards, and other devices are likely to take some configuring. Even manufacturer technical support tends not to be much help: they will tell users to simply try every combination until they find one that works, a trial-and-error approach that can be quite difficult and time-consuming, and may involve configuring jumpers and DIP switches as well as software. Many discs, too, are poorly documented, with only the briefest of installation and running instructions and no troubleshooting advice at all. Nonetheless, in general once the drive and disc are running they are very reliable and little further support is needed.

Information Sources

CD-ROM isn't always the best choice of access method for large databases, despite its large storage capacity. Many of the more expensive CD-ROM databases are also available in online versions, and it's important to take these into consideration in buying decisions. In general, CD-ROM makes sense if a large number of users need to access the same set of data. Many newspapers are now archived on CD-ROM, for example, and for regular or volume use the discs are a convenient and inexpensive way of setting up access.

But these discs are only updated quarterly, and individual years and publications are stored on separate discs. If a user needs to search over a longer time period or across publications, using separate discs rapidly becomes unwieldy. The online versions of the same databases are typically updated weekly or, in some cases, even daily, and most such services allow you to search across many hundreds of publications simultaneously. Such services aren't cheap - but neither is subscribing to a large number of disc sets. In addition, online you pay only for what you actually use, although at a higher rate.

Frequency of updates is even more of an issue with electronic telephone books, which are appealing in this day of directory assistance charges. Online versions of such databases are updated many times daily; the disc versions tend to be updated only quarterly or annually, not often enough for some businesses.

PCSA

The Author

Wendy M Grossman is a PC journalist and writer. She can be contacted as wendyg@cix.compulink.co.uk.

Understanding Microsoft At Work

Microsoft at Work, the company's blueprint for the office of the future, was announced last summer, but what's it all about? Jeffrey Goldberg explains.

Microsoft unveiled Microsoft At Work (MSAW) in June 1993 as a multi-company initiative to build office machines and computer products that proposed to offer:

- Ease of use, allowing users to access all product features, and to tailor the devices to their own preferences.
- A high-level of integration, allowing all devices in the workplace to communicate seamlessly with one another, in any combination.
- A strong platform, allowing Microsoft's partners to create a broad family of products and third parties to develop software.

The aim for Microsoft is to be the company selling the operating system and other software necessary to achieve these goals. This software will be both inside the product and in a connecting PC. The latter will, of course, work closely with the rest of the Windows family of products.

Microsoft At Work is so new a concept that few people outside Microsoft understand what it really is. This article, therefore, aims to explain the technology. We'll look at the types of products envisaged and benefits Microsoft sees from the incorporation of Microsoft At Work into these products.

Product Areas

Initially Microsoft believes the products that Microsoft At Work will burrow their way into are:

- Desktop and network-connected printers.
- Digital monochrome and colour copiers.
- Telephones and voice messaging systems.
- Fax machines and PC fax products.
- Handheld systems.

The first of these products is the Microsoft Windows Printing System, with the first range of non-Microsoft products likely to be networked fax machines.

Telephone System

Microsoft believes that, while the telephone is the most pervasive communications tool on the planet, it remains isolated from computer and other information devices. The company sees a new future for three separate sorts of devices:

- Microsoft At Work-based telephones. These will include desktop, public and portable phones, internally using the Microsoft At Work architecture and able to be integrated with PCs.
- MSAW-based PC phones. These will be Windows-based PCs which will allow users to tap into telephone networks connected via an add-in board, locally-connected telephone or a LAN.
- MSAW-based visual voice messaging servers. These are LAN-connected voice messaging systems that allow access to messages via Windows based PCs and MSAW-based telephones.

Microsoft believes the benefits of using Microsoft At Work for these products revolve around ease of use, visual voice messaging and integration with information systems.

Microsoft accepts that making a telephone call is about as simple as it is going to ever be, but it believes that the extra functionality provided by the network remains hidden from the user. They will overcome this using a graphical display, either on the phone itself or on a neighbouring Windows-based PC. This will add a menu based system that will present features in context and guide users through each step to accomplish a task. The system will also allow the users to prioritise and configure their phones to the way they work.

Answering machines and voice-mail have now become a ubiquitous part of the phone system yet Microsoft believes that today's systems are very inflexible. They point to the lack of checking for a messages from a particular person or being able to save a message for retrieval at some later date. The solution to this is to provide graphical management of voice messages which Microsoft calls visual voice messaging. This means that you can display voice messages as a list, much like an electronic mail inbox, including the caller's name or number and the time they called. The GUI would then have buttons to allow task such as forwarding the message to multiple people or returning the call.

This integration of a GUI, whether it be on the phone or in a PC, means that voice messaging can become part of a linked computer system. This could lead to many possibilities such as automatically showing reminders of things to discuss, or easier tracking of client calls for billing.

Microsoft At Work Fax

Microsoft believes that in spite of the dramatic growth in facsimile, it is a fairly primitive communications tool. They point to difficulty using machine features, low-quality printing, lack of integration with the work envi-

ronment and lack of security as some of the problems they intend to overcome with the MSAW products, of which there are four types:

- MSAW-based departmental fax machines. These are standalone, multi-functional machines which include a fax, printer, scanner and copier, that would be integrated completely with PCs and sharable via direct network or PC connection.
- MSAW-based fax servers. These are high volume, LAN-connected fax products.
- MSAW-based PC faxes. These are PCs with a standard modem capable of sending and receiving faxes, which will be integrated with the Windows environment via Microsoft at Work.
- MSAW Fax-Enabled Network. These are public networks that will be adding support for Microsoft At Work to allow users to broadcast facsimile messages in high volume and to access public mailboxes from any location.

Microsoft sees the main benefits of Microsoft At Work-based faxes as ease of use, the ability to send originals rather than copies, full document security, PC connectivity and MIS support. The ease of use comes from a graphical, touch-sensitive display which every Microsoft At Work fax machine will have. The ability to send originals comes from the ability to send and receive editable documents using a new fax protocol. This protocol can also use encryption to allow for more secure transmission than a normal fax machine. Microsoft, of course, would be sure to make the fax work with the rest of Windows including MAPI, the electronic mail programming interface.

"The integration with the rest of Windows makes fax machines a part of the MIS structure, which Microsoft says provides for easier management of costs."

The integration with the rest of Windows makes fax machines a part of the MIS structure, which Microsoft says provides for easier management of costs. Machines could use accounting codes to track costs, use off-peak rates or use the existing in-place data networks rather than public telephone networks. MIS could also ensure that they acted as a common resource sharing address books and sending the activity reports to accounts in binary form.

The MSAW Printer

At Microsoft, more than 25% of all calls received by the product support section involve printing problems of one type or another. It is natural, therefore, that part of Microsoft At Work targets printers to make them easier to install and use. This involves both hardware and software, but central to this is the addition of fast, bi-directional communications between printer and computer.

This would allow easier installation because the printer could then automatically load the correct driver. It would provide visual and audio feedback of both status and problems using a more intuitive interface than current models. Microsoft would also like the printer to use the same imaging technology as within Windows. This would guarantee WYSIWYG, where the type, graphics, and colours that appear on the screen are the same as those printed on the page. Microsoft also believes it will increase speed, avoiding the translation from the Windows imaging model to the printer's imaging language before printing requested documents. Finally, thinks Microsoft, it will result in a lower cost printer because the printer requires less processing power, less memory and fewer buttons.

"Microsoft sees the main benefits of Microsoft At Work-based faxes as ease of use, the ability to send originals rather than copies, full document security, PC connectivity and MIS support."

The MSAW Copier

The copier is a ubiquitous device in the office yet the use of analogue rather than digital technology leads Microsoft to believe it has a number of shortcomings:

- Since the print engine in the copier doesn't take digital input, users must print out documents before they can be copied.
- There is no way to integrate communications into today's copiers so that they function as document distribution tools. This is because they cannot receive, print, store or send image data in the same way a fax machine does today.
- Little special processing can be done on the documents beyond simple reduction and enlargement. For example, you cannot create numbered copies or stamp each page of each copy.

Microsoft At Work

"The OS is a completely new system that will support the real-time communications needs of office automation and telephony systems."

Microsoft assumes that this will change in the future leading to Microsoft At Work-based digital monochrome and colour copiers. They think these would be a leap over the current copiers being easier to use, having desktop connectivity, document distribution, true WYSIWYG colour and monochrome reproduction, image editing, on-demand reproduction and remote diagnostics.

In common with the other Microsoft At Work devices, the copier will have a touch-sensitive GUI to make it easier to use. The screen will allow previewing the document before printing and provide easy access to the other new imaging and output management functions. However, Microsoft wants the copier to be connected to a PC so that users can publish and distribute documents without any intermediate steps.

This publication would only be possible if the printing device produced output that looks the same on screen and on paper. This matching can only happen if both the computer and copier have the same page description and are colour calibrated to the same colour model. Microsoft plans to overcome this using the same imaging and colour models as used in Windows.

This will allow image editing on the device such as:

- Insertion of overlays to automatically number copies, stamp documents as confidential or add company logos.
- Colour accenting or other techniques will be available to emphasise a selected portion of the document.
- Optical character recognition filters (typically from third-parties) will convert analog text information into digital information.

Users will also be able to store these documents in the copier or on other LAN-connected servers. When they need a quick copy, they would just scroll through a list of documents, select the document and the number of copies. The copier being on the LAN, the administrator would be informed of any problem with the copier, and if the LAN was connected to a service centre it could allow technicians to run remote diagnostics on the copier.

MSAW Architecture

Microsoft has released little in-depth information about what the Microsoft At Work architecture is, apart from some top level detail. From what we have so far, MSAW appears to consist of:

- A new operating system "compatible" with the Windows API that is real-time, pre-emptive and multi-tasking.
- Communications software in both the Microsoft At Work machine and the desktop PC that will allow secure transmission of digital documents and the ability to control the machine.
- A rendering language to allow the transmission of documents with formatting and font intact.
- A simple graphical user interface in built into the machine and also an equivalent control panel on the controlling desktop PC.

We will go through the first three of these in further detail but little of this has been released, even in a Beta form, at the time of writing. All of the information therefore comes from open discussions with Microsoft and should be taken with a large dose of suspicion.

The Operating System

The OS is a completely new system that will support the real-time communications needs of office automation and telephony systems. It has the following key features:

- Pre-emptive, real-time support.
- Small memory footprint.
- Modular and device independent.
- Extensible.
- Windows-based, PC-compatible development environment using Windows compatible APIs.

MSAW Communications

Microsoft At Work communication services can use any available communication media such as serial ports, parallel ports, fax modems, data modems or LAN connections. It will allow any two users with a PC fax board, fax server or fax machine to exchange either editable or high-

"Communications software in both the Microsoft At Work machine and the desktop PC that will allow secure transmission of digital documents and the ability to control the machine."

quality "published" documents with the following features:

- Digital document transmission. This will come in one of two transmission types: High-quality "published" form using the Microsoft At Work rendering technology or as an editable document.
- **Security.** The Microsoft At Work message protocol employs public key/private key encryption, which will means that only the intended recipient can read the document when received. It also will allow users to request that the recipient be "authenticated" before delivering the message to be sure the recipient is the intended person. Finally, users will be able to include a "digital signature," which can be used to verify that the authenticity of the document.
- **Capabilities exchange.** When a Microsoft At Work based device connects to another Microsoft At Work-based device, the capabilities of those devices, such as what output resolutions are supported or whether the devices support colour, will be exchanged. This allows the sender to create the best possible format document for the recipient.
- **Compatibility with standard fax machines.** The message protocol is backward-compatible with existing fax machines using a variation of the Group III fax standard.
- **Integration with Windows messaging.** The Microsoft At Work message protocol interfaces with MAPI so that MAPI sees the At Work device as another type of recipient.

MSAW Rendering

Microsoft At Work Rendering is the model of the document inside the operating system. This model is usually different from the personal computer to the peripheral such as the printer or fax. In Microsoft At Work these models are the same. Microsoft believes this will allow At Work machines to:

- Harness the CPU and memory on the host to partition the work.
- Create a more efficient document printing and transmission format by eliminating redundancy and translation.
- Better match the screen representation to the output representation since they share the same model.

Conclusions

The difficulty of the Japanese entering the multi-functional market gives us some idea of the difficulty Microsoft will have entering the office equipment market. If the Japanese have had problems, then you can be sure that Microsoft won't find it easy.

The manufacturers and vendors in this market need to meet their customers requirements and fulfil their own profit margins. This doesn't always mean leading a technology race with another manufacturer since all the vendor wants is differentiation. The simple way is to offer product with better features for similar prices, for example:

- Printer companies add fonts, languages and resolution enhancement as improvements to their basic print engines.
- Copier companies add zoom, reduction and enlargement, and A3 copying.
- Facsimile manufacturers add auto-dialling and polling.

Unfortunately for the manufacturers, the number of new features over the past few years has dried up, leading to a search for a new style of product differentiation. This has led to an increasing reliance on the staples of lower price, better service and brand loyalty. If a promise of a radical technological change came along with the promise of differentiation nirvana, then as a manufacturer you'd look at it seriously. Microsoft wants the manufactures to believe this is Microsoft At Work.

Microsoft's public relations machine is one of the best in the business with pre-announcements and controlled leaks the norm. It is only when the product comes out that anyone can really assess the product, and by then the hype has built up to such an extent that it is sometime difficult to dislodge. We are at an early stage with Microsoft At Work, with a full launch but no product. We can't even obtain sufficient definitive information on the API or new operating system to come to any conclusions about fitness for purpose. In short, only Microsoft knows how strong Microsoft At Work really is.

It is, however, unusual that Microsoft chooses to launch the architecture with no real product demonstration. This smacks of desperation so we have to wonder why they did this. A cynic might answer that it is just another case of a big vendor spreading Fear, Uncertainty and Doubt to persuade customers from choosing a competitor's solution. There is no obvious competitor at present in all markets but there are competitors in each of the individual markets. By launching an all-encompassing strategy, Microsoft has trumped all its competitors who are present in individual markets. The trouble is: none of this actually exists in the shops.

In each of the markets mentioned Microsoft either hasn't competed in the market before or if it has, it has been unsuccessful. The bringing into Microsoft At Work of the Microsoft printing and fax strategy, smacks of FUD of the highest order. The addition of computer supported telephony brings Microsoft into an industry at a very late stage where the strategy will only work with the co-operation of major vendors, most are married.

In spite of this pessimism about the FUD factor, we have to admire the breadth of the strategy and the courage of attempting it. If ever there was an industry that needed a good shake-up then it's printing, copying and faxing.

PCSA

Acknowledgement

Extracted with permission from a report entitled "New Openings for Windows" by Jeffrey Goldberg, published by Dataquest.